

Launch Cloudbreak on AWS 2

Installing Cloudbreak on AWS

Date of Publish: 2019-05-28



<https://docs.hortonworks.com/>

Contents

Prerequisites on AWS.....	3
AWS account.....	3
AWS region.....	3
Virtual network.....	3
Security group.....	3
SSH key pair.....	3
Browser.....	4
Preparing the VM.....	4
System requirements.....	4
Root access.....	4
System updates.....	4
Iptables.....	4
Disable SELINUX.....	5
Docker.....	5
Configuring authentication with AWS.....	6
Create CloudbreakRole.....	6
Attach CloudbreakRole to the VM.....	10
Install Cloudbreak on a VM.....	11
Access Cloudbreak web UI.....	12
Next steps.....	14

Prerequisites on AWS

Before installing Cloudbreak, you must meet the following prerequisites:

AWS account

In order to install Cloudbreak on AWS, you must have an existing AWS account.

If you don't have an account, you can create one at <https://aws.amazon.com/>.

AWS region

Review available AWS regions and decide in which region you would like to launch the VM for Cloudbreak.



Note:

Clusters created via Cloudbreak can be in the same or different region as Cloudbreak; When you launch a cluster, you select the region in which to launch it.

Related Information

[Regions and Availability Zones \(AWS\)](#)

Virtual network

You must have a virtual network configured on your cloud provider.

Security group

Ports 22 (SSH), 80 (HTTPS), and 443 (HTTPS) must be open on the security group.

SSH key pair

In order to access the Cloudbreak VM via SSH, you will be required to use your SSH key pair.

If you do not have an SSH key, import an existing key pair or generate a new key pair in the AWS region in which you are planning to launch the VM for Cloudbreak.

Steps

1. Navigate to the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Check the region listed in the top right corner to make sure that you are in the correct region.
3. In the left pane, find NETWORK AND SECURITY and click Key Pairs.
4. Do one of the following:
 - Click Create Key Pair to create a new key pair. Your private key file will be automatically downloaded onto your computer. Make sure to save it in a secure location. You will need it to SSH to the cluster nodes. You may want to change access settings for the file using `chmod 400 my-key-pair.pem`.
 - Click Import Key Pair to upload an existing public key and then select it and click Import. Make sure that you have access to its corresponding private key.

You need this SSH key pair to SSH to the Cloudbreak instance and start Cloudbreak.

Related Information

[Amazon EC2 key pairs \(AWS\)](#)

Browser

In order to access Cloudbreak web UI, you should use one of the following supported browsers: Chrome, Firefox, or Safari.

Preparing the VM

To install the Cloudbreak deployer and install the Cloudbreak application, you must have an existing VM.

You should launch the VM by using the steps provided in your cloud provider documentation. Once you have the VM ready, review the following requirements:

System requirements

In order to install Cloudbreak, your system must meet the minimum requirements.

Ensure that your system meets the following requirements:

- Minimum VM requirements: 16GB RAM, 40GB disk, 4 cores
- Supported operating systems: RHEL, CentOS, and Oracle Linux 7 (64-bit)



Note:

Default prewarmed and base images provided with Cloudbreak for cluster creation cannot be used for installing Cloudbreak. Instead, you can use an image provider or cloud provider default image with one of the operating systems supported by Cloudbreak.



Note:

You can install Cloudbreak on Mac OS X for evaluation purposes only. Mac OS X is not supported for a production deployment of Cloudbreak.

Root access

Every command mentioned in this documentation must be executed as root.

In order to get root privileges execute:

```
sudo -i
```

System updates

Perform these steps to ensure that your system is up-to-date.

To ensure that your system is up-to-date, run:

```
yum -y update
```

Reboot the VM if necessary.

Iptables

Perform these steps to install and configure iptables.

Steps

1. Install iptables-services:

```
yum -y install net-tools ntp wget lsof unzip tar iptables-services
systemctl enable ntpd && systemctl start ntpd
systemctl disable firewalld && systemctl stop firewalld
```

**Note:**

Without iptables-services installed the iptables save command will not be available.

2. Configure permissive iptables on your machine:

```
iptables --flush INPUT && \
iptables --flush FORWARD && \
service iptables save
```

Disable SELINUX

Perform these steps to disable SELINUX.

Steps

1. Disable SELINUX:

```
setenforce 0
sed -i 's/SELINUX=enforcing/SELINUX=disabled/g' /etc/selinux/config
```

2. Run the following command to ensure that SELinux is not turned on afterwards:

```
getenforce
```

3. The command should return “Disabled”.

Docker

Perform these steps to install Docker.

The minimum Docker version is 1.13.1. If you are using an older image that comes with an older Docker version, upgrade Docker to 1.13.1 or newer.

Steps

1. Install Docker service:

CentOS 7

```
yum install -y docker
systemctl start docker
systemctl enable docker
```

RHEL 7

```
yum install yum-utils
yum-config-manager --enable rhui-REGION-rhel-server-extras
yum install -y docker
systemctl start docker
systemctl enable docker
```

2. Check the Docker Logging Driver configuration:

```
docker info | grep "Logging Driver"
```

3. If it is set to Logging Driver: journald, you must set it to “json-file” instead. To do that:

a. Open the docker file for editing:

```
vi /etc/sysconfig/docker
```

b. Edit the following part of the file so that it looks like below (showing log-driver=json-file):

```
# Modify these options if you want to change the way the docker daemon
runs
OPTIONS='--selinux-enabled --log-driver=json-file --signature-
verification=false'
```

c. Restart Docker:

```
systemctl restart docker
systemctl status docker
```

Configuring authentication with AWS

Before you can start using Cloudbreak for provisioning clusters, you must select a way for Cloudbreak to authenticate with your AWS account and create resources on your behalf.

There are two ways to do this:

- **Key-based:** This is a simpler option which does not require additional configuration at this point. It requires that you provide your AWS access key and secret key pair in the Cloudbreak web UI later.
- **Role-based:** This requires that you or your AWS admin create an IAM role (referred to as the CloudbreakRole) to allow Cloudbreak to assume AWS roles (the AssumeRole policy). Later, you must create another IAM role (referred to as the CredentialRole) for Cloudbreak to be able to perform specific actions via Cloudbreak credential.

If you chose the key-based option, you do not need to do anything at this point and can proceed to the next step.

If you chose the role-based option, proceed to create the CloudbreakRole and then attach the role to your VM. For more information on how role-based authentication works, refer to [Authentication with AWS](#). For instructions on how to create and attach the CloudbreakRole, refer to the documentation linked below:

Related Information

[Create CredentialRole](#)

Create CloudbreakRole


If using role-based authentication, perform these steps to create the CloudbreakRole.

Use the following AssumeRole policy definition:


```
{
  "Version": "2012-10-17",
  "Statement": {
    "Sid": "Stmt1400068149000",
    "Effect": "Allow",
    "Action": [ "sts:AssumeRole" ],
    "Resource": "*"
  }
}
```

Steps

1. Navigate to the IAM console > Roles and click Create Role:



2. In the “Create Role” wizard, select AWS service role type and then select any service:



3. When done, click Next: Permissions to navigate to the next page in the wizard.
4. Click Create policy and the create policy wizard will open in a new browser tab:

Create role

1 Trust 2 Permissions 3 Review

Attach permissions policy

Choose one or more policies to attach to your new role. Each role can have a default maximum of 10 attached policies.

Create policy Refresh

Filter: Policy type Search Showing 287 results

	Policy name	Type	Attachments	Description
<input type="checkbox"/>	AdministratorAccess	Job function	3	Provides full access to AWS services and resources.
<input type="checkbox"/>	AmazonAPIGatewayAd...	AWS managed	0	Provides full access to create/edit/delete APIs in Amazon A...
<input type="checkbox"/>	AmazonAPIGatewayInv...	AWS managed	0	Provides full access to invoke APIs in Amazon API Gateway.
<input type="checkbox"/>	AmazonAPIGatewayPus...	AWS managed	0	Allows API Gateway to push logs to user's account.
<input type="checkbox"/>	AmazonAppStreamFullA...	AWS managed	0	Provides full access to Amazon AppStream via the AWS M...

- 5. Select the JSON view, and then copy and paste the policy definition. You can either copy it from the section preceding these steps or download and copy it from [here](#).

Create policy

1 2

A policy defines the AWS permissions that you can assign to a user, group, or role. You can create and edit a policy in the visual editor and using JSON. [Learn more](#)

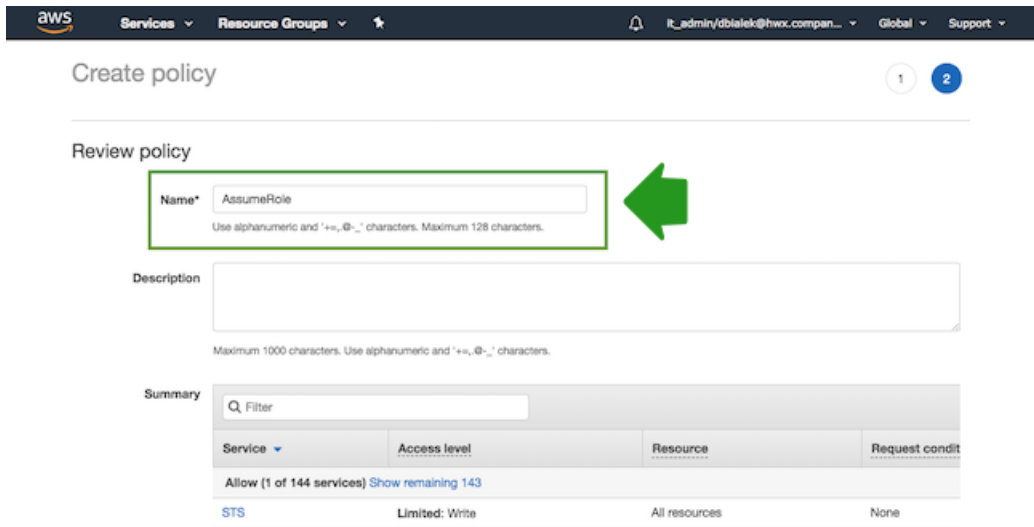
This policy validation failed and might have errors converting to JSON : The policy must have at least one statement For more information about the IAM policy grammar, see [AWS IAM Policies](#)

Visual editor **JSON** Import managed policy

```
1- {
2-   "Version": "2012-10-17",
3-   "Statement": {
4-     "Sid": "Stm1400068149000",
5-     "Effect": "Allow",
6-     "Action": ["sts:AssumeRole"],
7-     "Resource": "*"
8-   }
9- }
```

- 6. When done, navigate to Review policy.

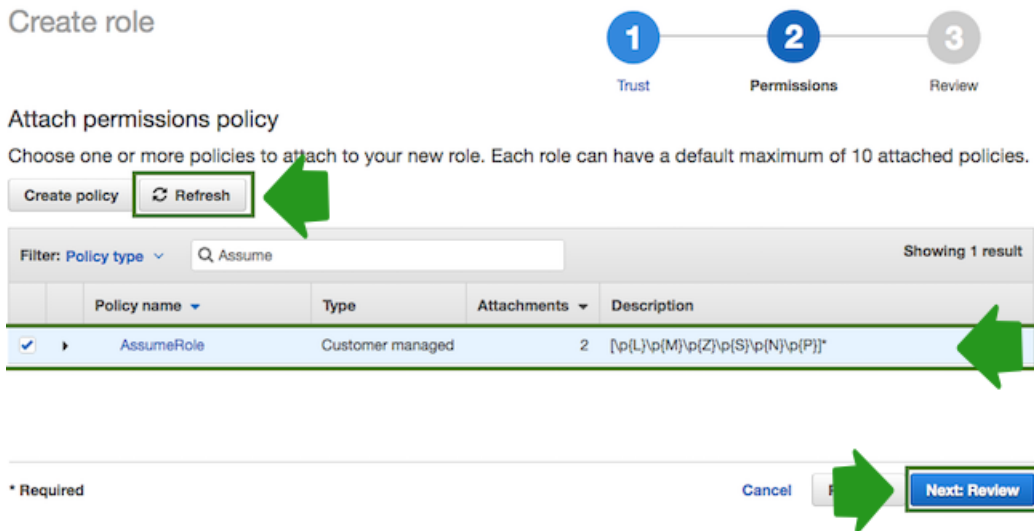
7. On the Review policy page, in the Name field, enter a name for your policy, such as “AssumeRole”:



8. When done, click Create Policy.

9. Return to the previous browser tab where you started creating a new role (since the create policy wizard was opened in a new browser tab).

10. Click Refresh. Next, find the policy that you just created and select it by checking the box:



11. When done, click Next: Review.


12. In the Roles name field, enter role name, for example “CloudbreakRole”:

Create role

1 Trust 2 Permissions 3 Review

Review

Provide the required information below and review this role before you create it.

Role name* 


Maximum 64 characters. Use alphanumeric and '+=,@-.' characters.

Role description

Maximum 1000 characters. Use alphanumeric and '+=,@-.' characters.

Trusted entities The identity provider(s) cloudformation.amazonaws.com

Policies [AssumeRole](#)

* Required Cancel Previous Create role 

13. When done, click Create role to finish the role creation process.

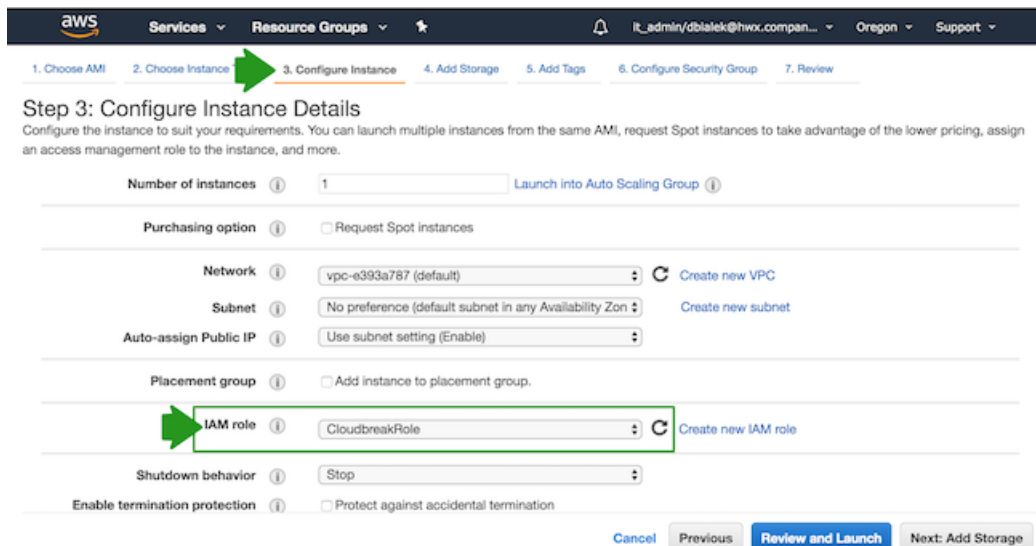
As an outcome of this step, the CloudbreakRole should be created in your IAM console on AWS.

Attach CloudbreakRole to the VM


If using role-based authentication, after creating the CloudbreakRole, attach it to your VM. You can do this during the VM launch process or once your VM is running.

Attach CloudbreakRole when launching the VM

The option to attach an IAM role is available from the instance launch wizard Step 3: Configure Instance:



aws Services Resource Groups IT_admin/dbialek@hwx.compan... Oregon Support

1. Choose AMI 2. Choose Instance  3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 3: Configure Instance Details

Configure the instance to suit your requirements. You can launch multiple instances from the same AMI, request Spot instances to take advantage of the lower pricing, assign an access management role to the instance, and more.

Number of instances [Launch into Auto Scaling Group](#)


Purchasing option Request Spot instances

Network [Create new VPC](#)

Subnet [Create new subnet](#)

Auto-assign Public IP

Placement group Add instance to placement group.

 IAM role [Create new IAM role](#)

Shutdown behavior

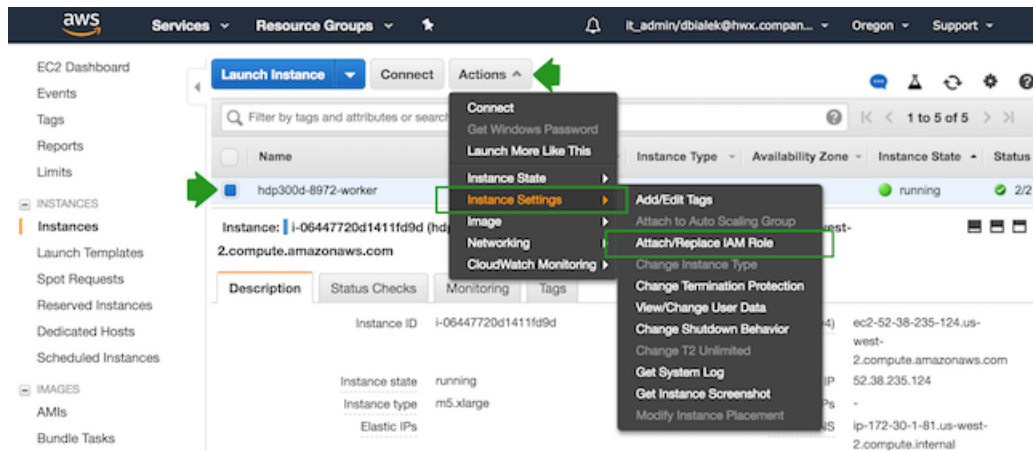
Enable termination protection Protect against accidental termination

Cancel Previous Review and Launch Next: Add Storage

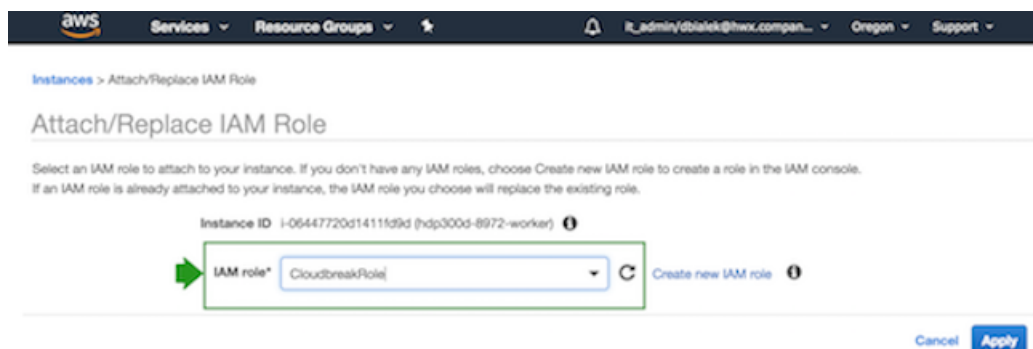
Attach CloudbreakRole to an existing VM

Perform the following steps from the EC2 console on AWS:

1. Select the EC2 instance that you are planning to use for Cloudbreak and from the Actions menu select Instance Settings > Attach/Replace IAM Role:



2. Next, under IAM role, select the CloudbreakRole that you created earlier and click Apply to apply the configuration:



As an outcome of this step, the CloudbreakRole IAM role should be attached to the VM on which Cloudbreak will be launched on AWS.

Install Cloudbreak on a VM

Install Cloudbreak on your own VM from a Cloudbreak deployer binary.



Attention: As of December 31, 2021, Cloudbreak reached end of support. For more information, see [Support lifecycle policy](#). Cloudera recommends that you migrate your workloads to [CDP Public Cloud](#).

Steps

1. Install the Cloudbreak deployer and unzip the platform-specific single binary to your PATH. For example:

```
yum -y install unzip tar
curl -Ls public-repo-1.hortonworks.com/HDP/cloudbreak/cloudbreak-deployer_2.9.1_$(uname)_x86_64.tgz | sudo tar -xz -C /bin cbd
cbd --version
```

Once the Cloudbreak deployer is installed, you can set up the Cloudbreak application.

2. Create a Cloudbreak deployment directory and navigate to it:

```
mkdir cloudbreak-deployment
cd cloudbreak-deployment
```

3. In the directory, create a file called Profile with the following content:

```
export UAA_DEFAULT_SECRET=[ $MY-SECRET ]
export UAA_DEFAULT_USER_PW=[ $MY-PASSWORD ]
export UAA_DEFAULT_USER_EMAIL=[ $MY-EMAIL ]
export PUBLIC_IP=[ $MY_VM_IP ]
```

For example:

```
export UAA_DEFAULT_SECRET=MySecret123
export UAA_DEFAULT_USER_PW=MySecurePassword123
export UAA_DEFAULT_USER_EMAIL=test@cloudera.com
export PUBLIC_IP=172.26.231.100
```

You will need to provide the email and password when logging in to the Cloudbreak web UI and when using the Cloudbreak CLI. The secret will be used by Cloudbreak for authentication.

You should set the `CLOUDBREAK_SMTP_SENDER_USERNAME` variable to the username you use to authenticate to your SMTP server. You should set the `CLOUDBREAK_SMTP_SENDER_PASSWORD` variable to the password you use to authenticate to your SMTP server.

4. Generate configurations by executing:

```
rm *.yml
cbd generate
```

The `cbd start` command includes the `cbd generate` command which applies the following steps:

- Creates the `docker-compose.yml` file, which describes the configuration of all the Docker containers required for the Cloudbreak deployment.
- Creates the `uaa.yml` file, which holds the configuration of the identity server used to authenticate users with Cloudbreak.

5. Start the Cloudbreak application by using the following commands:

```
cbd pull-parallel
cbd start
```

This will start the Docker containers and initialize the application. The first time you start the Cloudbreak application, the process will take longer than usual due to the download of all the necessary docker images.

6. Next, check Cloudbreak application logs:

```
cbd logs cloudbreak
```

You should see a message like this in the log: Started Cloudbreak Application in 36.823 seconds. Cloudbreak normally takes less than a minute to start.

Related Information

[Troubleshooting Cloudbreak](#)

Access Cloudbreak web UI

Log in to the Cloudbreak web UI by using the following steps.

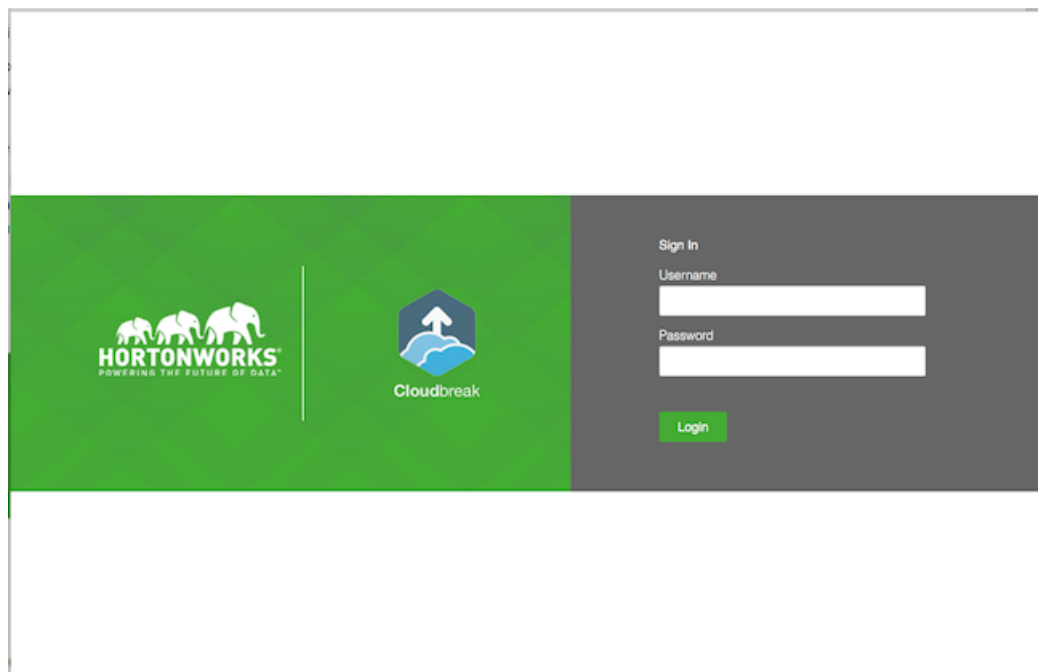
Steps

1. You can log into the Cloudbreak application at https://IP_Address. For example https://34.212.141.253. You may use cbd start to obtain the login information. Alternatively, you can obtain the VM's IP address from your cloud provider console.
2. Confirm the security exception to proceed to the Cloudbreak web UI.

The first time you access Cloudbreak web UI, Cloudbreak automatically generates a self-signed certificate, due to which your browser warns you about an untrusted connection and asks you to confirm a security exception.

Browser	Steps
Firefox	Click Advanced > Click Add Exception... > Click Confirm Security Exception
Safari	Click Continue
Chrome	Click Advanced > Click Proceed...

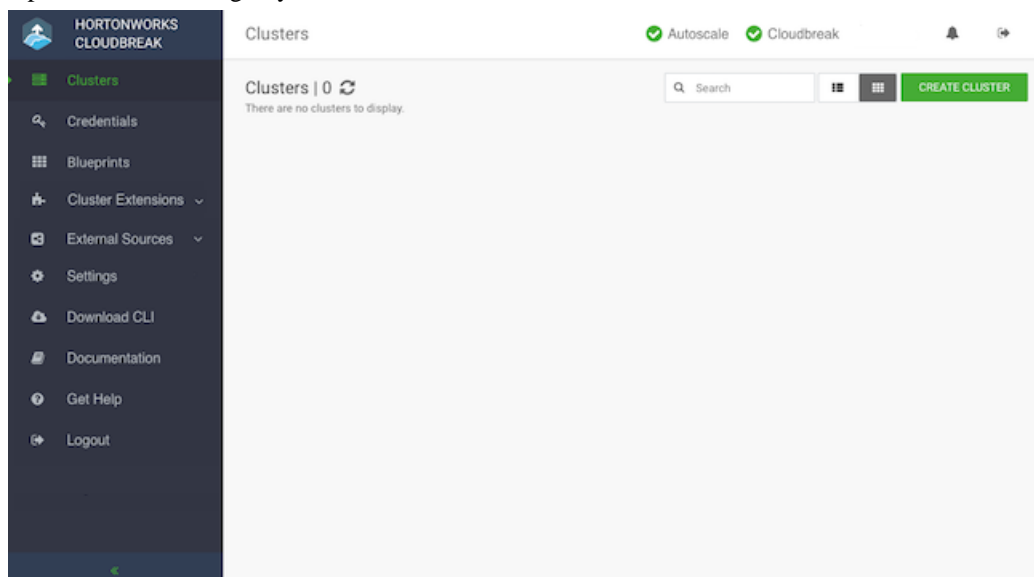
3. The login page is displayed:



4. Log in to the Cloudbreak web UI using the credentials that you configured in your Profile file:

- The username is the UAA_DEFAULT_USER_EMAIL
- The password is the UAA_DEFAULT_USER_PW

5. Upon a successful login, you are redirected to the dashboard:



Next steps

After launching Cloudbreak, you must configure an external Cloudbreak database (if using Cloudbreak for production) and then create a Cloudbreak credential.

Configuring an external Cloudbreak database

By default, Cloudbreak, uses an embedded PostgreSQL database to persist data related to Cloudbreak, configuration and so on. This database is only suitable for non-production Cloudbreak, deployments. For production, you must configure [an external Cloudbreak database](#).

Creating a Cloudbreak credential

This step is required. Only after you've performed it, you can start creating clusters. There are two ways to create a Cloudbreak credential on AWS. If you are just getting started with Cloudbreak, we recommend using the easier key-based credential method. For Cloudbreak credential options on AWS, refer to [Credential options on AWS](#).

Related Information

[External Cloudbreak database](#)

[Cloudbreak credential options on AWS](#)

[Configure Cloudbreak](#)